

I. Introduction

In the fall of 2008, when the U.S. auto industry approached a state of crisis and at least General Motors and Chrysler faced a serious chance of bankruptcy, one proposal explored was a merger between the two firms. Such a proposal—a merger of the largest and fourth largest companies in an already moderately concentrated market—would have seemed outlandish only a few years earlier, due in part to the likely disapproval by the Department of Justice, but now was seriously considered, with little or no public discussion of the possible anti-competitive effects. This paper attempts to model the effects of such a merger by answering the following questions:

1. What are the plausible new pricing and product line scenarios that GM might pursue after merging with Chrysler? That is, which models would GM continue to sell, which would they cut due to cannibalization, and could they increase profits by raising (or lowering) prices?
2. How would consumer surplus change between the pre- and post-merger scenarios?

The answers to these questions are crucial to the consideration of policy regarding the American car industry, especially as policy-makers try to balance interests of the shareholders and employees of the companies against the potentially lost consumer welfare resulting from a merger.

To answer my primary questions, I build a model of the U.S. retail light vehicle market. Because the products are differentiated and consumers typically make infrequent

purchases of a quantity of just 1, I model demand using a discrete-choice, differentiated-products framework. A rich, relatively recent literature models demand in car industry this way. In the model, each consumer's utility from a given vehicle is a function of several consumer characteristics, vehicle characteristics, parameters (coefficients), and a random error term that contains all the unobserved influences. The consumer chooses the vehicle that gives him/her the highest utility. Given certain assumptions about the distribution of the error term, the econometrician can calculate a probability for each vehicle that it is a particular individual's utility-maximizing choice. Summing a vehicle's probability of being purchased across all the new car buyers in the U.S. gives the vehicle's market share. I specifically follow Train and Winston's 2007 paper, using their demand model and their estimates of the parameters that govern it. I update it by calibrating the vehicle dummies to closely approximate true, 2009 market shares of vehicles.

The demand model simulates the market shares of the vehicles in the choice set that consumers face, under certain conditions. I run simulations to model the effects of a merged GM-Chrysler dropping vehicles, as well as raising prices. I find that while there are several good candidates for vehicles that the firm could consider dropping, the market competition prevents them from profitably raising prices much at all, and in fact consumer welfare is only negligibly affected by the merger, under certain scenarios.

II. Industry and literature review

2.1 Industry review

The U.S. retail light vehicle market comprises a large portion of consumer spending. Though sales have recently fallen significantly—down 18% from 16.1 million in 2007 to 13.2 million in 2008 (Visnic, 2009)—in 2008, motor vehicles and parts contributed \$379.9 billion, or 2.66%, to the United States GDP, according to data from the Bureau of Economic Analysis. The industry is also relatively concentrated. According to data from Automotive News, in the first quarter of 2009, the “Big 6” leading automakers—GM, Toyota, Ford, Chrysler, Honda, and Nissan—accounted for 82.5% of vehicles sold in the U.S. The Herfindahl-Hirschman Index (HHI) for 2008 was 1324, which meets the Department of Justice criterion for a “moderately concentrated” industry. Figure 1 shows the market shares of the Big 6.

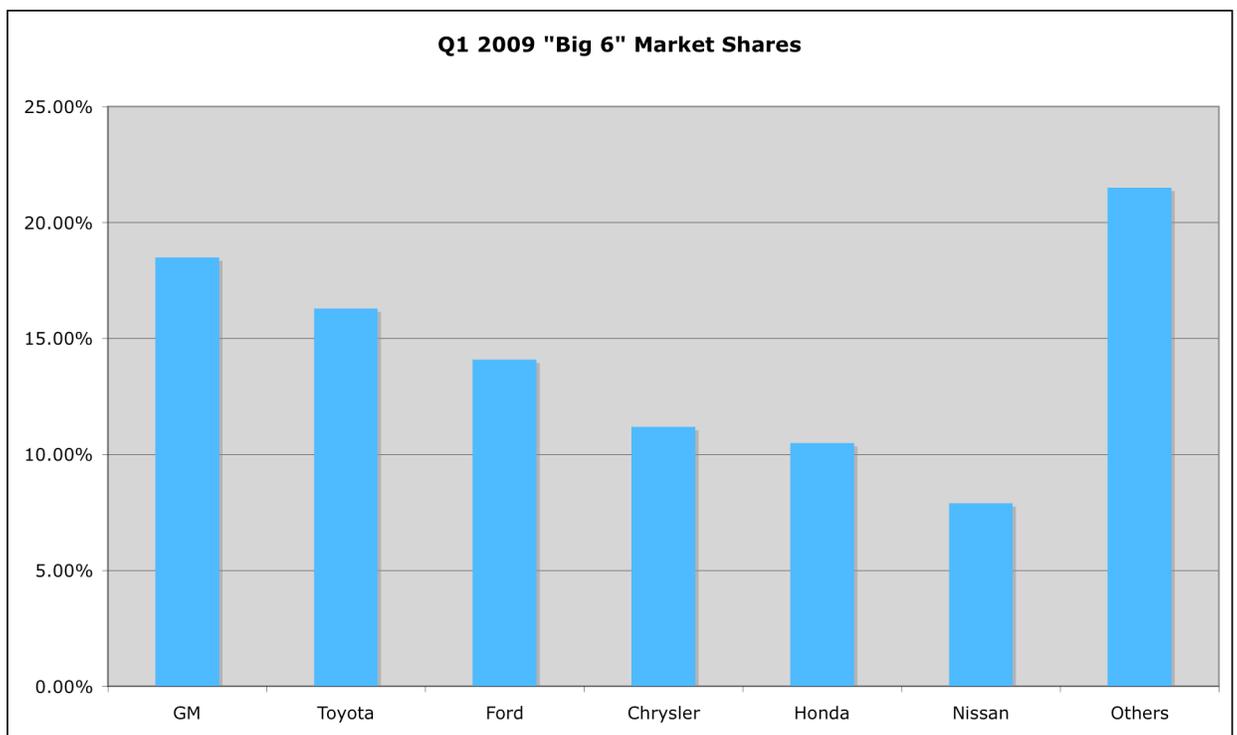


Fig. 1: Market shares of the firms with the six highest market shares in the U.S.

The “Big 3” domestic automakers have seen their collective market share fall steeply in the last several years, from 71% in 1999 to 50% in 2008. The speed of this

drop, however, masks the long road that brought the domestics to this point. American car companies shared in the postwar boom of the 1950s and 60s, serving a growing country while insulated from any significant competition. The oil crisis and appearance of cheaper, more fuel-efficient Japanese imports changed all that in the 1970s, bringing particular pain to an over-extended, liquidity-constrained Chrysler. In 1979, Chrysler teetered on the verge of bankruptcy, and was only spared by emergency loan guarantees of \$1.5B from the federal government. The bailout proved successful: Chrysler actually repaid the loans early, thanks to the rollout of new front-wheel drive vehicles as part of Lee Iacocca's turnaround efforts. Meanwhile, Congress moved to protect the industry from Japanese imports by erecting trade barriers, including tariffs and "voluntary" export restrictions. Under these auspices, the U.S. car companies were able to temporarily regain their dominance, and in fact made great headway with the introduction of new vehicles such as the minivan and SUV, both of which Chrysler was crucial to piloting with its Dodge Caravan and Jeep lines.

The Japanese and other foreign competitors responded by investing heavily in production capacity on American soil, allowing them to eventually escape much of the burden of the trade barriers. A recent Wall Street Journal editorial points out that by the end of 2008, "foreign" automakers employed 113,000 Americans, compared to 239,000 working for U.S. carmakers ("America's Other Auto Industry"). The workers at these "transplant" factories are mostly not unionized, and have significantly lower wages and benefits than the unionized employees of the Big 3. Toyota, Honda, and Nissan also continued to improve the quality and image of their mostly smaller cars, a segment in which American automakers were happy to cede market share while they reaped the huge

profit margins on SUVs and pickup trucks. Years of lax fuel economy legislation permitted the car companies to forego the expensive task of making these larger vehicles fuel-efficient.

The price of gasoline climbed throughout the 2000s, finally reaching a critical height in 2007 that reversed the trend toward more pickup and SUV purchases. Then, as the country slipped into recession in 2008, car sales dropped precipitously for all the major companies, with GM and Chrysler losing the most thanks to the even greater collapse of SUV and pickup sales. At the same time, the American automakers' long-growing pension and healthcare liabilities were catching up with them, so that their costs remained high even as revenues plummeted. News stories reported that both companies were using up their cash reserves quickly, and the credit crisis that erupted in the fall made it all the more difficult for the companies to raise cash. Ford, GM, and Chrysler appealed to Congress for the extension of \$25B in loans that had been approved, but not yet disbursed, for the development of more fuel-efficient vehicles. The loans were eventually granted, beginning an escalating process of government involvement and support for the Big 3.

In October, word leaked of merger talks between GM and Chrysler. Some of the principal motivations discussed in the press were access to liquidity, and cutting costs. Specifically, it was suggested that GM could trade some of their 51% share of GMAC (their financing arm) to the private equity group Cerberus, which owns Chrysler as well as the other 49% of GMAC, in exchange for additional cash. The acquisition process would also give GM the opportunity to re-negotiate labor contracts with Chrysler employees, cut underperforming brands, and hopefully exploit "cost synergies," by

sharing some of the redundant fixed costs of production across the expanded company. Bill Vlasic and Andrew Ross Sorkin discussed these possibilities in an October 11, 2008 *New York Times* article, and also voiced the concern of this paper that such a merger could face opposition from the Antitrust Division of the Department of Justice. However, on October 28, Vlasic and Edmund Andrews reported in *The New York Times* that the Bush administration was considering aid to help support the merger.

The merger talks fell through, and the idea was officially tabled by the start of 2009. Today, in mid-April 2009, the future of the two companies remains very much in doubt. The Obama administration has extended significant lifelines to the companies and begun taking an active role in their management. To date, Obama has removed the CEO of GM, Rick Wagoner, and directed Chrysler to form a partnership with the Italian automaker Fiat by June. The possibility of bankruptcy proceedings for one or both companies looms larger than ever, as the Obama administration has publicly discussed a “surgical bankruptcy,” to permit the companies restructure and escape much of their debt. Fritz Henderson, GM’s new CEO, has resisted the idea, instead offering equity swaps to debt-holders, but acknowledges that if significant concessions are not agreed upon, bankruptcy may be unavoidable. The idea of a merger has been re-opened, as observers have predicted that Chrysler may not survive as an independent entity. Neil Winton reported on April 3, 2009 in the *Detroit News* about a private research report modeling industry profits in the event that Chrysler is taken over and pared down by GM, which it also predicts would shed several lines of its own, such as Saturn and Pontiac. Rob Cox and Antonie Curry argued in the *New York Times* on April 14, 2009 that a GM-Chrysler merger remains the best option for the two companies’ viability. While it appears less

likely than it did in the fall of 2008, a merger in the next several months cannot be ruled out.

2.2 Literature review

My thesis draws on two major strains of research in applied microeconomics. The first is a series of papers exploring the use of discrete-choice demand models in differentiated-products industries, and several of these focus on the auto industry in particular. Bresnahan (1987) pioneered their application to the auto industry. Goldberg (1995) and Berry, Levinsohn, and Pakes (1995) built substantially on the model, working in different ways to advance the econometric techniques for estimating the demand parameters: Goldberg focused on a nested logit structure to better organize levels of consumer choice, while Berry *et al.* tried to allow for much more flexible substitution patterns, and pioneered a class of estimation techniques now widely known by the first initials of the three authors, BLP. In 2002, Amil Petrin showed how discrete-choice demand models could be put to use estimating changes to consumer welfare, as he did with the case of the minivan—I make similar estimates of consumer surplus changes here. In the most recent paper along these lines, Train and Winston (2007) augment the typical Consumer Expenditure Survey data with their own, privately conducted survey, which asked new-car buyers what their second choice would have been, allowing for a much more nuanced estimation of the demand parameters.

For my analysis of a merger, I draw on a second strain of literature exploring the usage of merger simulations in differentiated-products industries, used in both academic and applied antitrust settings. Werden *et al.* (1996) survey the usage and drawbacks of

merger simulations based on logit demand models, so-called because market shares are derived from a ratio of logit formulas (all of the papers mentioned earlier rely on some variant of a logit model). In his 2000 paper, Nevo simulates a merger in the ready-to-eat cereal industry and estimates changes to consumer welfare. One goal of this paper is to make use of the econometric developments in these diverse fields, applying them specifically to the U.S. auto industry.

III. Theory and Simulation design

3.1 Supply

The structure of supply in the retail car market in the United States is best characterized as an oligopoly. Automakers are assumed to maximize profit by their choice of product lines and prices. Because this paper is concerned with GM's actions in the short run—what prices might they change and what vehicles might they drop soon after a merger—it makes sense to ignore the more long-term firm decisions of which vehicles to design and create, and with which attributes.

Competing in a somewhat concentrated market does not mean, *prima facie*, that a firm has the ability to set prices above marginal cost and extract greater consumer surplus. Bertrand's Paradox demonstrates that, assuming no collusion, firms in a *very* concentrated industry may end up with little to no profits: in his example, a duopoly of two firms who strategically choose prices on identical (perfectly substitutable) goods, each with ample production capacity and the same marginal cost, will undercut each other until they reach an equilibrium of pricing at marginal cost—the same outcome as perfect competition.

The car market's key difference is product differentiation. No two vehicles are perfect substitutes: they are differentiated by price, numerous specifications like weight and horsepower, the availability of options such as navigation systems, and intangibles such as reputation and brand effects. Consider the extreme branding case of the Chevy Silverado and GMC Sierra pickup trucks: they are the same truck, with identical specifications and MSRP, both built by GM, but simply badged with different nameplates. There were 67,283 Silverados and 22,508 Sierras sold in the first quarter of 2009. Thus if we imagine vehicles competing in a multi-dimensional characteristic space, where "distance" (in a topological sense) apart represents substitutability between goods, no two products would occupy the same point. In an extreme sense, one could say each manufacturer has a monopoly in the market for each particular model it sells. But of course product differentiation does not end competition between products; it only dampens it. GM knows if it cut the price of the Sierra by 10% or improved its features with no change in price, some consumers would substitute to it from other models, including from the Silverado. In order to model firms' strategic, profit-maximizing choices of prices and product lines, we must have a model for the behavior of consumers in response to suppliers' choices.

3.2 Demand

Consumers in the U.S. retail light vehicle market are households (this paper ignores fleet sales to corporate customers). I assume that the consumer maximizes her utility by her choice of vehicle. I restrict my attention to those consumers who have chosen to buy a new vehicle. This is equivalent to the strong assumption that consumers don't change their decision of *whether* to buy a new car in response to the pricing and

attributes of new vehicles—they only change their choice of which new car to buy.

Consumers typically only buy one vehicle at a time. Therefore, the market-level quantity demanded for a particular vehicle is really the number of people buying that vehicle—one person is not expected to buy more units of a vehicle model if its price falls, but is rather *more likely* to buy one unit of a vehicle model. For this reason, we need a model of demand that gives a probabilistic account of the discrete choices that consumers make.

The first step is to model the consumer's utility function. As mentioned in Section 2, there is an entire literature modeling demand in the car market; here, I specifically follow Train and Winston (2007). Consumers derive utility from the attributes of a vehicle, and their preferences are assumed to be heterogeneous. That is, different consumers derive different levels of utility from a vehicle conditional on certain attributes of the individuals: for instance, a person with higher income will derive less disutility from a vehicle's cost than a person with a lower income. The utility a consumer n derives from a vehicle j is given by

$$U_{nj} = \delta_j + \beta' X_{nj} + \mu_n' w_{nj} + \varepsilon_{nj} \quad (1)$$

The first two terms capture what might be called the systematic portion of utility. The δ_j term represents the utility that does not vary across consumers—for instance, how much people, on average, value different levels of the power-to-weight ratio in a car. X_{nj} is a matrix of vehicle attributes interacted with consumer attributes, like price and income, and it is multiplied by a vector β of coefficients. The 3rd and 4th terms represent the idiosyncratic portion of utility. The term w_{nj} is another matrix of consumer and vehicle attributes, but the coefficients μ_n are allowed to vary by consumer, in order to capture

random, unobserved tastes that are correlated. This random-coefficients portion is the defining characteristic of mixed-logit models (for a discussion, see Train, 2003, ch. 6).

The term ε_{nj} is a random error term assumed to be distributed independently and identically according to the extreme value distribution. The error term includes all unobserved contributors to utility.

If all of utility could be accounted for by the systematic, observed portion of the utility function above, the researcher could know with certainty the utility that every individual gets from every vehicle, provided he has data on the relevant attributes of the consumers and vehicles, as well as knowledge of the demand parameters. Market demand for each vehicle would simply be the total of number consumers for whom that vehicle provides the greatest utility.

In reality, the econometrician faces incomplete information about the components of utility. However, due to the assumption that the error terms are distributed i.i.d. extreme value, Train and Winston (2007) show that the probability P_{ni} that a consumer n buys a vehicle i is given by the logit formula

$$P_{ni} = \int \frac{e^{\delta_i + \beta' X_{ni} + \mu_n' w_{ni}}}{\sum_j e^{\delta_j + \beta' X_{nj} + \mu_n' w_{nj}}} f(\mu) d\mu \quad (2)$$

where the summation is across all the vehicles j that the consumer is facing. The integration over the distribution of μ integrates out the random coefficients, weighting the logit ratios by the conditional probability distribution function of μ . This enables the

substitution pattern to escape the infamous IIA (independence of irrelevant alternatives) consequence of the logit formula.

These probabilities provide the link between individual vehicle choices and the aggregate demand. The market share of a particular vehicle is the sum, over all new-car buyers, of the probability that each buyer purchases it.

When considering how well this model accounts for real-world consumer behavior, it is clear that most consumers probably do not actually research the prices and specifications of every vehicle model and trim package on the market. Rather, they make a series of decisions on particularly important aspects of a vehicle, narrowing their search until they only consider a handful of specific models in great detail. In Train and Winston's (2007) survey of recent new car-buyers, the most common answer by far to the question "How many other vehicles did you seriously consider aside from the one purchased?" was "one." From the econometrician's standpoint, the nested decisions by which consumers reach their vehicle choices are particularly important because the nests are closely tied to substitution patterns. I follow the literature (particularly Goldberg (1995)) in assuming there are five important decision levels, as illustrated in figure 2: whether to buy a car at all, whether to buy it new or used, vehicle segment, vehicle origin, and particular model.

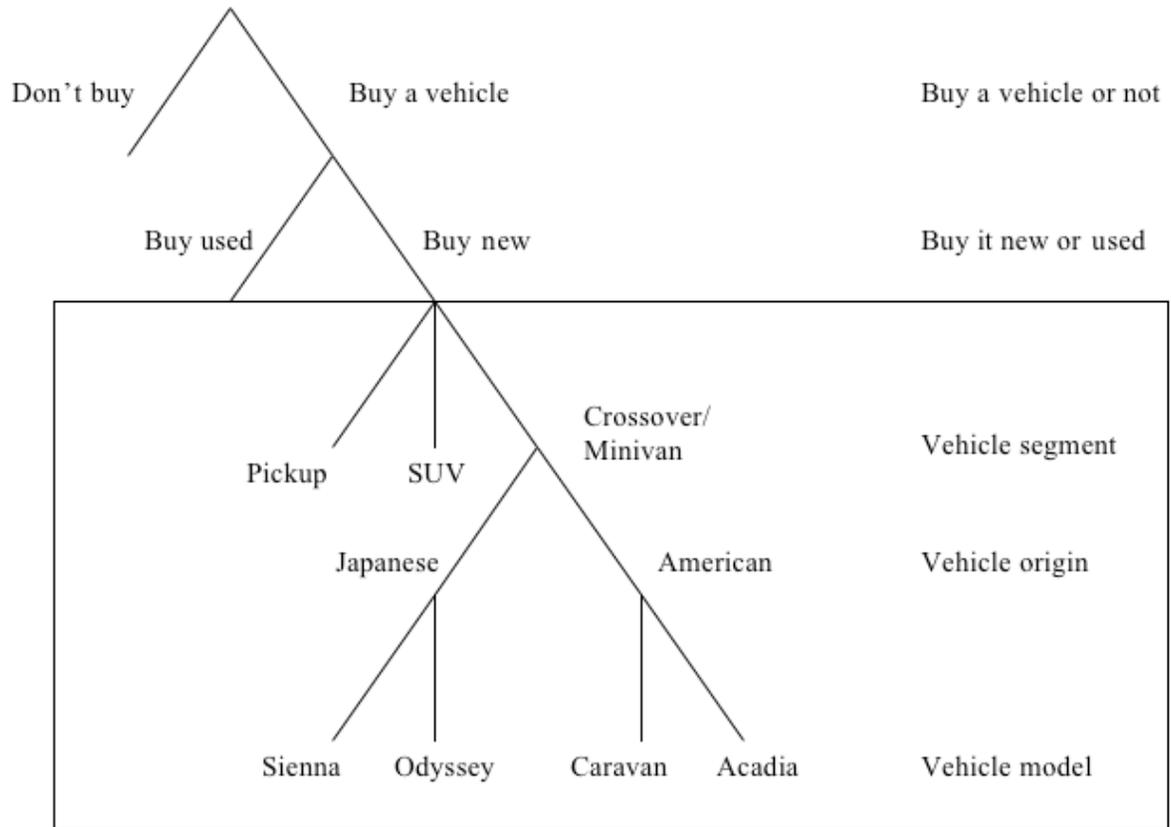


Fig. 2: Decision tree for a household. Note that I consider only the choices in the boxed portion of the decision tree, so that all the decisions are conditional on having chosen to buy a new car. Also, only a few representative choices are shown here at each of the three decision levels within the box.

People begin to narrow their field of choices with the question of vehicle segment—do they want to buy a pickup truck, mid-sized sedan, etc.? It seems reasonable to put this decision highest on the decision tree because the differences in nature of the vehicles among segments are the most fundamental, and so the difference in the level of inter-nest substitution (eg. a consumer being torn between a compact car and a pickup truck) and the level of intra-nest substitution will be the greatest of any of the decisions.

I assume that the next level of the decision tree is the vehicle manufacturer's geographic origin, reflecting the way some consumers will tend to compare mostly among Japanese imports, or mostly among American cars. There is strong evidence in

publications like Consumer Reports that consumers attribute certain reputational characteristics to all of the manufacturers from particular countries; for instance, the conventional wisdom holds that Japanese car companies produce the most reliable vehicles, and American companies make somewhat less reliable vehicles. Other preferences for certain nationalities of cars vary more by consumer: some Americans consider it a moral choice, supporting domestic jobs, to “Buy American,” and thus presumably only consider the American-made cars in their vehicle segment of choice. Nonetheless, it is important to note that a considerable amount of inter-nest substitution is likely at this level. It seems intuitively reasonable that someone considering the Honda Accord would most likely be comparing it to the Toyota Camry, but it is also possible that he or she would compare it to the Chevy Malibu.

Finally, the consumer chooses a particular vehicle in his preferred nationality’s subset of his preferred segment. His choice here likely has some level of sorting as well, to arrive at the final two or three candidates to fully research, but that decision is so close to the person’s specific utility function that it would be useless to try to generalize. The choice will be driven by the vehicle’s traits, the consumer’s traits, and the consumer’s particular tastes, as they relate to the utility they get from a vehicle.

3.3 Merger effects

The nested nature of the consumer decision is crucial to my analysis of the effects of a merger between GM and Chrysler. In this framework, firms are constrained from exercising too much market power by the availability of competing substitutes. The efficacy of a substitute is determined by how close it is to a given product in characteristics, but this nesting scheme suggests that a good way to identify candidates

for close substitutes is to look at similar cars in the same nests. A merger would permit the merged firm, which I will call GM-Chrysler for the sake of simplicity, to increase market power and harm consumer welfare in one of two ways. First, GM could identify segments in which a GM and a Chrysler model were each other's closest substitutes, and then eliminate the worse selling of the two, while raising the other's price.

By combining the largest and third-largest of the three American car companies, this merger would be expected to most augment GM's market power in a segment where the American car companies are most dominant. As it happens, the American automakers' strength in a segment increases with the size of vehicles in it: they are least successful in compacts, the most dominant in truck-based ("real") SUVs and pickup trucks, and reasonably competitive in the minivan and cross-over segment. Thus, pairs of GM and Chrysler vehicles in these latter three segments are the greatest cases for concern about market power, and the starting point for analysis in the empirical section.

Firms' pricing and vehicle line decisions are generally assumed to be in a Bertrand-Nash oligopoly equilibrium. Ideally, I would solve for the best-response decisions comprising the new Bertrand-Nash equilibrium under my different simulated counterfactuals, but the dimensions of such a solution—countless discrete combinations of vehicle line possibilities, each of which would need its own equilibrium solution for the prices that several multi-product firms set—place it outside the scope of my thesis. Instead, I fix the short-run pricing decisions of the rest of the firms in the industry for my counterfactuals.

3.4 Empirical design

In order to simulate the market shares of vehicles and manufacturers under new equilibria in the U.S. car industry, I use Train and Winston's (2007) estimates of the demand parameters. Train and Winston (2007) report that consumer-level survey data, with information on choice rankings, proved necessary for the accurate estimation of parameters. Because I lack such data, and due to non-trivial computational requirements for estimation in any case, I choose to simply rely on these previously estimated parameters. I use data on a sample of 6,477 households from the Consumer Expenditure Survey and 55 of the best-selling 2009 model-year vehicles. I include in my sample of households those whose most recently purchased car was bought new, a very similar requirement to Train and Winston's conditioning trait of having purchased a new 2000-model-year car.

The behavior of the simulations depends very much on how I choose which vehicles to consider. I choose not to include *all* 2009-model-year cars and light trucks because they number over 200, and due to the distribution of sales—most sales are concentrated in a few popular models—there are rapidly diminishing returns to including more than roughly 50 of them. At the same time, the computational requirements of the simulations rise roughly with the square of the number of models. I construct my vehicle choice set in terms of market segments: pickup trucks, SUVs, crossovers/minivans, mid-sized sedans, and compact cars. I exclude luxury and full-sized cars for three reasons: they make up only a small portion of sales in the market, so their model-by-model sales would be out of proportion to the sales of vehicles in the five major segments given above; they are not very close substitutes for vehicles in these five segments, being generally much more expensive; and finally, Chrysler has virtually no presence in the

luxury and full-sized segments, which means that consumers forced to substitute away from Chrysler models, due to higher prices or their discontinuation, are especially unlikely to shift into luxury and full-sized cars. To choose the vehicles in each segment, I begin by including their most-sold new 2009 model-year vehicles, and then sequentially add in the next-best-selling models, such that each segment sample's share of sales among these five samples closely reflects the share of all the vehicles in that segment relative to all the sales of the five segments. Figure 3, below, illustrates:

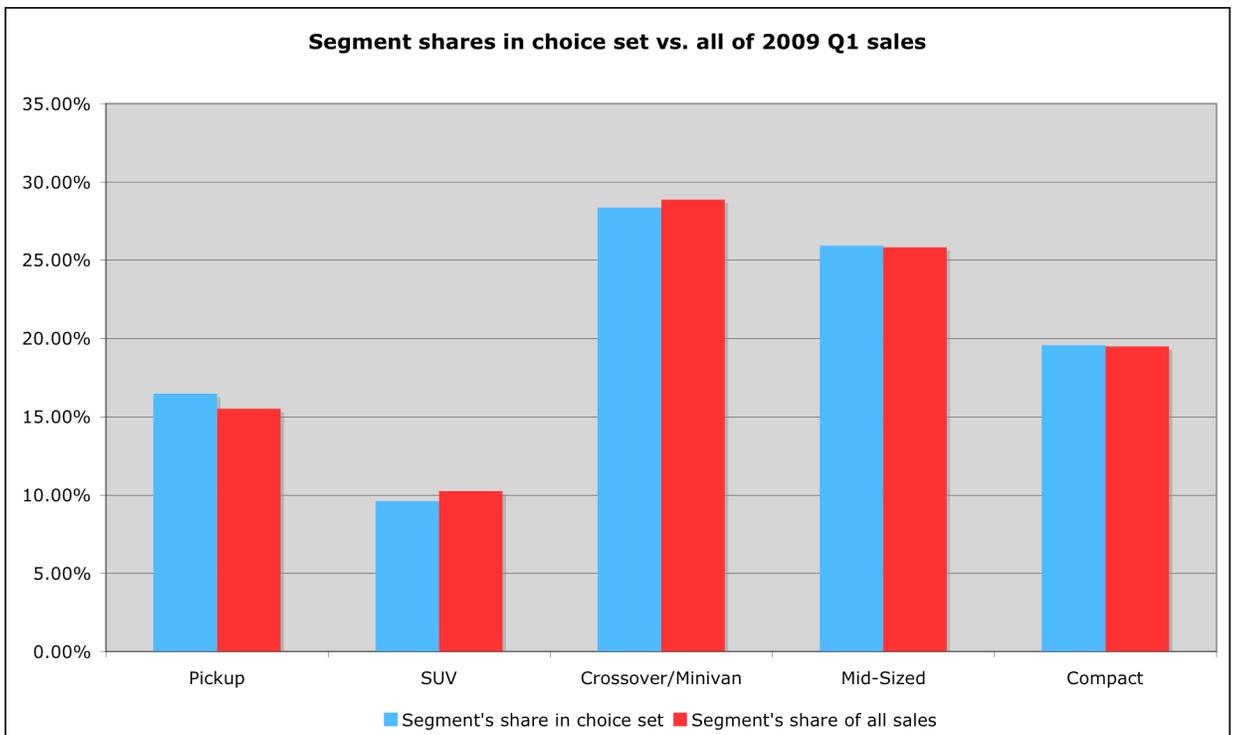


Fig. 3: Market shares by segment. The blue bar represents the share of the total sales of vehicles in the choice set contained in that segment. The red bar represents the share of all sales in the market accounted for by that segment.

Note that I choose to fill my choice set by segment, rather than trying to match the market shares of manufacturers in the choice set to their shares in the overall market.

Nonetheless, **figure 4** shows they are mostly quite close:

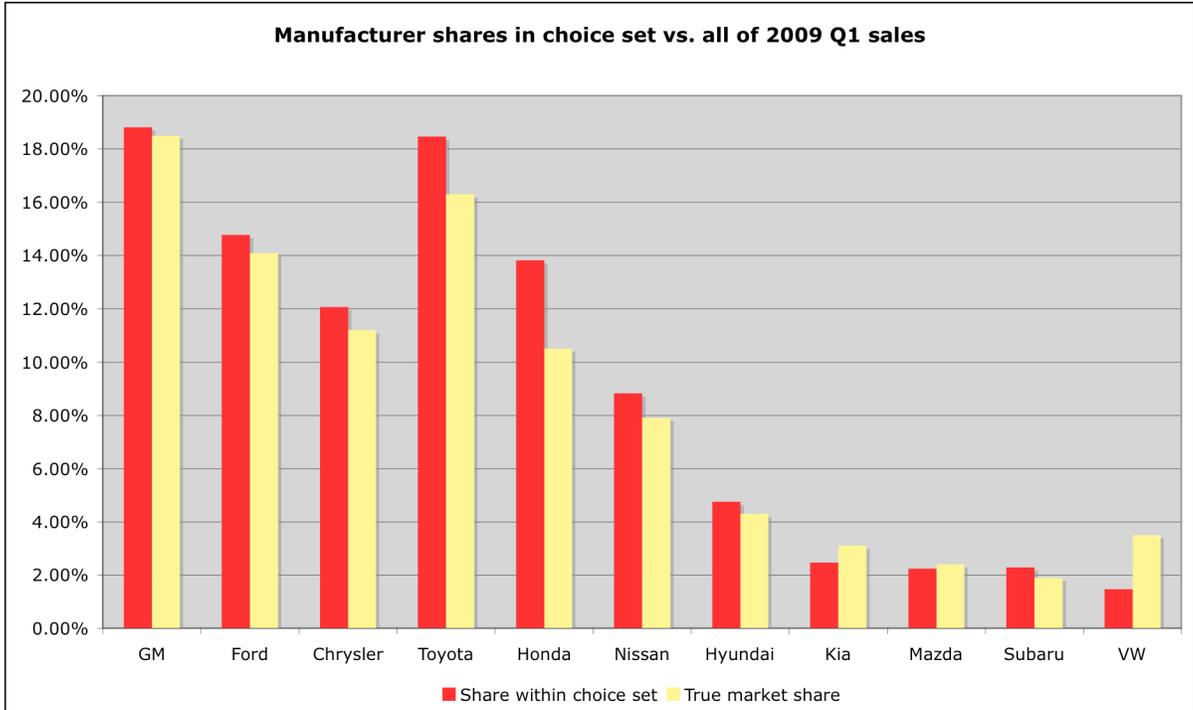


Fig. 4: The market shares of each manufacturer. Each red bar shows that manufacturer’s share of sales among vehicles in the choice set, and the yellow bar shows that manufacturer’s true market share.

The Japanese manufacturers exhibit significant disparities. This is because, relative to the American manufacturers, they tend to sell fewer models under fewer brands, and therefore their sales are concentrated in a few bestsellers. The Accord and Civic alone, for instance, comprised 46% of Honda’s sales in Q1 2009, whereas the F Series and Escape accounted for 35% of Ford’s sales. Because I focus on the best-selling several vehicles in each segment, it is not surprising that the Japanese are over-represented relative to the Americans. While regrettable, this aspect of the choice set could only be mitigated, somewhat, by including many more vehicles in each segment and would only be fully solved with the inclusion of every vehicle—each of Honda and Toyota, in fact, has only two crossovers/minivans, one mid-sized car, and one compact car included in the choice set, plus one truck for Toyota, so cutting down on their number of models is

certainly not the solution. In any case, the primary distortion this could introduce to the simulations would be to *under*-predict the degree to which consumers consider GM and Chrysler vehicles as close substitutes, and thus under-predict GM's increased market power following a merger.

I follow Train and Winston's choice of consumer and vehicle traits. The consumer variables included are after-tax income as calculated by the CES, age, sex, urban/rural designation, the presence of children under the age of 16, and the make of the household's most recent car purchase (which must have been purchased new, by the requirements for being included in the sample). The vehicle attributes I gather data on are the manufacturer's suggested retail price, horsepower, curb weight, wheelbase, end-to-end length, EPA-rated gallons per mile (translated to early-2000s numbers to match with Train and Winston's parameters), the Consumer Reports expected reliability rating, and the Kelley Blue Book 2-year retained value prediction. All variables except the last two are courtesy of Automotives News; the last two are from their respective websites. The following tables provide summary statistics on my sample of consumers, and compare them to those with which Train and Winston estimated their model:

Socioeconomic variables	My sample	Train & Winston's sample
Average after-tax income	\$73,083.51	\$67,767
Average age	52	54
Percent male	50.11%	54%
Percent with child aged 1-16	25.23%	19%
Percent who live in rural location	5.34%	45%

Table 1: Characteristics of heads of household in my sample vs. Train and Winston (2007).

Manufacturer origin of most recently bought car	My sample	Train & Winston's sample
	Proportion	Proportion
U.S.	55.7%	64%
Japanese	36.2%	28%
European	4.2%	5%
Other	4.0%	3%

Table 2: The distribution of households' most-recently purchased vehicles, by national origin. Note that my sample is in line with the trend toward fewer American cars, as it was drawn in 2005 and 2006.

To run a simulation, I begin by calculating the predicted utility \hat{U}_{nj} given to each consumer n from each vehicle j given by equation 1 exclusive of the error term. I then compute the probability, P_{nj} given by equation 2; to approximate integration over the distribution of μ_{nj} , which is given to be standard normal, I take the simple mean over 200 random draws from the distribution, which I save and re-use for all simulations to minimize error. I also weight the different consumers' probabilities by their CES population weights. This step helps to correct for aspects of these consumers that are not representative of the general population—for instance, the over-representation of urban dwellers. But it suffers the drawback of counter-weighting away some of those unrepresentative aspects of this sample that were caused by conditioning on having recently purchased a new car—for instance, above-average income. The simulated market share for a given vehicle j is then simply the sum of the weighted probabilities P_{nj} over all the consumers in my sample, divided by the number of consumers.

Average utility	Coefficient
Constant	-6.8520
Manufacturer's suggested retail price (<i>in thousands of dollars</i>)	-0.1063

Expected retained value after 2 years (<i>in thousands of dollars</i>)	0.0550
Horsepower divided by weight (<i>in tons</i>)	0.0312
Automatic transmission dummy	0.6787
Wheelbase (<i>inches</i>)	0.0509
Length minus wheelbase (<i>inches</i>)	0.0279
Fuel consumption (<i>in gallons per mile, times 10⁴ for scaling</i>)	-0.0032
Luxury or sports car dummy	-0.0558
SUV or station wagon dummy	0.7231
Minivan dummy	-1.1288
Pickup truck dummy	0.0661
Chrysler dummy	0.0654
Ford dummy	0.2696
General motors dummy	0.3715
European dummy	2.4008
Korean dummy	0.8017

Utility that varies over observed consumer characteristics	Coefficient
Manufacturer's suggested retail price divided by household income	-1.6025
Vehicle reliability based on the <i>Consumer Reports</i> repair index (1-5, 1 being worst) for women aged 30 and over (0 otherwise)	0.3949
Luxury or sports car dummy for lessors (1 if vehicle is luxury or sports car and consumer leased last car, 0 otherwise)	0.6778

Minivan dummy for households with an adolescent (1 if vehicle is a van and household has childgren aged 7 to 16, 0 otherwise)	3.2337
SUV or station wagon dummy for households with an adolescent	2.0420
Number of previous consecutive GM purchases	0.3724
Number of previous consecutive GM purchases for households in a rural location	0.3304
Number of previous consecutive Ford purchases	1.1822
Number of previous consecutive Chrysler purchases	0.9652
Number of previous consecutive Japanese manufacturer purchases	0.7560
Number of previous consecutive European manufacturer purchases	1.7252

Utility that varies over consumers unrelated to observed characteristics (random coefficients portion)	Coefficient
MSRP divided by respondent's income, times a random standard normal	0.8602
Horsepower times a random standard normal (divided by 10^4 for scaling)	45.06
Fuel consumption (<i>gallons per mile, times 10^4 for scaling</i>) times a random standard normal	-0.0102
Light truck, van or pickup dummy, times a random standard normal	6.8505

Table 3: Table of coefficients estimated by Train and Winston (2007) that I use. I omit standard errors because they did not affect my analysis—I used all the point estimates of the coefficients as they appear in this table.

The only parameter estimates not published in Train and Winston's paper are the vehicle dummy coefficients. These parameters are important for capturing unobserved, average effects of people's preferences for particular vehicle models. The dummies were last estimated for 2000-model-year vehicles, making them dubious at best; many of the vehicles in the 2000 choice set no longer exist (e.g. Ford Taurus), several new models have been introduced (including many in the rapidly growing crossover segment), and

Each red bar shows the market share predicted by a simulation with the baseline parameters, after calibration of vehicle dummies.

I compute cross-price elasticities numerically rather than analytically, by raising the price of a particular vehicle j in the data by 5%, and running a new simulation. Then the cross-price elasticity of demand for vehicle i with respect to the price of vehicle j is given by $(\Delta Q_i / \Delta P_j) * (P_j^0 / Q_i^0)$. I do so for all of GM and Chrysler's vehicles, as a guide to some of the substitution patterns in the demand system.

Model	Silverado	Ram	Wrangler	Equinox	Traverse	Malibu	Impala	Cobalt	wncountry	Liberty	Tahoe	Cherokee	Journey	Acadia	Caravan
Silverado	-2.299	0.175	0.053	0.020	0.063	0.034	0.023	0.015	0.031	0.040	0.045	0.020	0.018	0.052	0.007
Ram	0.278	-2.575	0.073	0.016	0.049	0.031	0.021	0.012	0.044	0.056	0.038	0.028	0.024	0.039	0.011
Wrangler	0.222	0.195	-2.572	0.018	0.057	0.028	0.020	0.011	0.094	0.071	0.046	0.037	0.030	0.048	0.030
Equinox	0.239	0.122	0.052	-2.911	0.078	0.020	0.013	0.007	0.055	0.044	0.054	0.022	0.020	0.064	0.012
Traverse	0.227	0.112	0.049	0.024	-3.252	0.013	0.009	0.005	0.058	0.044	0.053	0.024	0.022	0.077	0.009
Malibu	0.078	0.045	0.015	0.004	0.008	-2.972	0.138	0.048	0.008	0.009	0.012	0.004	0.003	0.006	0.008
Impala	0.069	0.039	0.014	0.003	0.007	0.180	-3.218	0.048	0.007	0.008	0.010	0.004	0.002	0.006	0.008
Cobalt	0.119	0.060	0.021	0.005	0.010	0.163	0.124	-2.562	0.009	0.011	0.013	0.005	0.004	0.008	0.009
Towncountry	0.138	0.125	0.100	0.021	0.071	0.015	0.011	0.005	-2.955	0.079	0.053	0.043	0.035	0.061	0.034
Liberty	0.207	0.186	0.087	0.019	0.063	0.020	0.014	0.008	0.091	-2.760	0.046	0.036	0.030	0.053	0.023
Tahoe	0.219	0.117	0.054	0.022	0.072	0.025	0.017	0.008	0.058	0.043	-4.124	0.023	0.018	0.059	0.015
Cherokee	0.200	0.178	0.089	0.019	0.066	0.019	0.014	0.007	0.097	0.070	0.047	-3.489	0.031	0.057	0.024
Journey	0.216	0.182	0.085	0.020	0.073	0.014	0.010	0.006	0.093	0.070	0.044	0.037	-2.545	0.063	0.018
Acadia	0.229	0.111	0.051	0.024	0.095	0.012	0.009	0.005	0.060	0.045	0.053	0.025	0.023	-3.490	0.009
Caravan	0.110	0.113	0.115	0.016	0.041	0.056	0.043	0.018	0.119	0.070	0.048	0.038	0.023	0.034	-2.803
G6	0.076	0.043	0.015	0.004	0.008	0.179	0.139	0.051	0.007	0.008	0.011	0.004	0.003	0.006	0.008
Vue	0.245	0.119	0.052	0.024	0.087	0.015	0.010	0.006	0.057	0.045	0.053	0.024	0.022	0.074	0.010
Patriot	0.159	0.103	0.040	0.017	0.053	0.016	0.010	0.005	0.043	0.036	0.043	0.017	0.014	0.041	0.010
Yukon	0.218	0.116	0.053	0.022	0.072	0.025	0.017	0.008	0.058	0.043	0.058	0.023	0.018	0.059	0.015
Suburban	0.224	0.115	0.055	0.023	0.078	0.021	0.015	0.007	0.061	0.045	0.057	0.025	0.020	0.066	0.014

Table 4: Selected cross-price elasticities. Each entry is the percent change in quantity demanded of the row vehicle, given a 1% increase in the price of the column vehicle. The Caravan column illustrates some interesting substitution patterns for former Caravan buyers.

Table 4 shows some cross-price elasticities of GM and Chrysler vehicles. The Caravan column is highlighted to illustrate some of the interesting information it contains about substitution patterns. It shows that the Town & Country is its closest substitute, unsurprisingly, because that is an almost identical minivan, but it also shows that the Jeep Wrangler, Liberty, and Cherokee are the next three closest substitutes, even though they are SUVs. This suggests that those individuals with a significant probability of buying a Dodge Caravan are doing so largely out of manufacturer loyalty, and will therefore

switch to some remaining vehicle made by Chrysler if they get priced out of their first choice.

IV. Results

4.1—Choice of vehicles to drop

After merging, General Motors has eight Chrysler vehicles in the choice set from which it will likely choose to drop some. In order to predict the merged firm's choice of vehicle to drop, it does not make sense to simply choose the vehicle with the highest cross-price elasticities with the rest of GM's models. As Table 4 shows, and as intuition would confirm, the Dodge Ram, being Chrysler's best-selling vehicle, has by far the highest cross-price elasticities with other GM models. This is a consequence of the Ram's high market share—when its price is raised 1%, and 2.575% of its consumers substitute away from it, those 2.575% of its former buyers are much more numerous than the 3.489% of Cherokee consumers who would be driven away by a 1% price increase, because the Dodge Ram began with roughly 8 times as many buyers. Therefore as they substitute to other vehicles, the former Dodge Ram buyers cause similarly greater increases in the quantity demanded.

GM's choices of vehicles to drop are motivated by profit maximization. Unfortunately, industrial cost functions are notoriously opaque to the econometrician. However, it is widely agreed-upon that the merged GM-Chrysler would choose to cut at least a few product lines due to the fixed costs associated with each one. Therefore, if we assume that there are similar fixed cost savings to be had for all models, then a first approximation of profit-maximization is revenue-maximization.

After dropping one vehicle, the merged firm's revenue would fall due to the direct loss of revenue from the sales of that car, but that drop would be somewhat offset as some proportion of the former buyers of the dropped vehicle substitute to other products in the GM-Chrysler line. This accounts for why the Dodge Ram, in spite of its high cross-price elasticities noted above, would not be the first choice of the merged firm to drop, because its first-order direct revenue losses would be too great.

One other assumption I make when solving for the merged firm's choice of vehicles to drop is that GM would be much more likely to drop Chrysler models, and in fact I assume they only choose among Chrysler vehicles to drop. The first reason for this is that, during the takeover process, GM would likely have greater flexibility for closing Chrysler plants, because of the likelihood that the deal would be brokered either by the Obama administration or a bankruptcy court. In either case, the emergency nature of the proceedings would likely grant GM some protection from the costly aspects of plant closings, such as employee buyouts, and other labor-related costs. Secondly, due to the costs of incorporating new production processes, management teams, etc., every plant of Chrysler's that GM decides to keep running will add an additional merger-related fixed cost, which could be entirely foregone by simply never taking the plant over. Third, the management of GM will almost certainly control most of the process (except to the degree that the government interferes, perhaps), and it is likely that they would prefer to cut other brands simply out of having greater loyalty or familiarity for their own product lines. Although an extraordinary streamlining of the final company's product offerings is possible, in which case some GM lines might be cut as well, for all of the above reasons

it seems fair to assume that the *first* products chosen to be cut in any scenario would be at least a few Chrysler vehicles.

To approximate the merged firm’s revenue-maximizing choice, I compute the change in GM-Chrysler’s total revenue caused by the loss of any single Chrysler vehicle. To do this, I run eight simulations, dropping one Chrysler vehicle from the choice set each time, and then compute the change in total revenue. The results appear below:

Vehicle dropped:	Change in total revenue (in millions)
Dodge Ram	-\$5,696.96
Jeep Wrangler	-\$1,862.23
Chrysler Town & Country	-\$2,114.81
Jeep Liberty	-\$1,669.01
Jeep Grand Cherokee	-\$1,011.15
Dodge Journey	-\$646.01
Dodge Caravan	-\$560.45
Jeep Patriot	-\$1,211.61
Starting revenue:	\$81,447.50

Table 5: This table shows by how much annual revenue is calculated to fall under my model of demand. Note that the model predicts market shares of every vehicle, not sales—to arrive at sales, I assume 13 million vehicles sold annually and multiply by the predicted market shares. Revenue is given by sales*MSRP.

The bottom four vehicles—the Grand Cherokee, Journey, Caravan, and Patriot—have significantly smaller negative impacts on total revenue than the other four. I assume these are the four vehicles most likely to be dropped.

These choices are also in accord with the hypothesis that GM-Chrysler would try to avoid cannibalization of sales between two vehicles in a given segment. The Dodge

Caravan is a very similar minivan to the Chrysler Town & Country, so the merged firm might decide it only needs one of the two. Similarly, the Jeep Patriot and Jeep Grand Cherokee are significant competitors with the Jeep Wrangler, Jeep Liberty, and GM's Chevy Tahoe, so their elimination helps to consolidate the SUV segment. Finally, the Dodge Journey, Chrysler's new crossover offering, is a poorly-selling offering in the already crowded segment, where GM has the GMC Acadia, Chevy Equinox, and Saturn Vue.

4.2—Price changes

I hypothesized that the merged firm would attempt to exploit its consolidated market power by raising prices. It is a straightforward expectation from oligopoly theory that as the number of firms competing approaches 1, prices tend to rise. I am not able to solve for the new Bertrand-Nash oligopoly pricing equilibrium after a merger, so instead I consider a unilateral move by GM-Chrysler, to raise prices 5% across the board. While the equilibrium would almost certainly not feature such a uniform price change, due to the different price elasticities of the different goods, I believe this move is reasonable, because it captures some of the consumer surplus of those people motivated strongly by manufacturer loyalty. The cross-price elasticities table suggests that many Chrysler buyers' preferences are driven a great deal by brand loyalty, because Chrysler products in radically different segments, such as the pickup and the crossover, exhibit unusually high cross-price elasticities.

The results, however, suggest that a unilateral, across-the-board price raise would not increase GM-Chrysler's revenue:

Scenario:	Total annual revenue (in millions):
All original vehicles, original prices	\$81,447.50
Dropped four vehicles above, original prices	\$77,950.37
Dropped four vehicles above, raised all prices 5%	\$72,765.96

Table 6: The merged firm's total annual revenue, assuming 13 million total vehicles sales for the entire U.S. car market. A unilateral price hike would not be profitable for GM-Chrysler, even after cutting several former Chrysler models.

Table 6 shows that GM-Chrysler could not profitably raise its prices 5%, unilaterally.

4.3—Consumer surplus effects

With a unilateral price hike out of the question, the conservative prediction for post-merger behavior simply remains that GM-Chrysler would stop selling the four vehicles mentioned above, due to cost considerations. The consumer surplus effects of such a move have two parts: the loss of all of the consumer surplus on the discontinued models, and the gain in consumer surplus from the shifting-out of the demand curves of every remaining vehicle. The change is unambiguously expected to be negative, because every buyer of the discontinued vehicles is now being forced to buy their second or third choice, which gives them lower utility.

To compute consumer surplus, I take a linear approximation of every demand curve, close to the point (Q^*, P^*) of the market quantity and price. For each demand curve, consumer surplus is computed as the area under the curve, above the market price. I compute once for the baseline scenario, and once for the scenario where the Grand Cherokee, Caravan, Journey, and Patriot have been dropped from the choice set. Every demand curve shifts out, and possibly rotates, after the removal of the four vehicles, because the forced substitution from those vehicles results in greater demand for any

other vehicle at every price. Table 7 summarizes the changes in computed consumer surplus:

Consumer surplus before dropping vehicles (millions):	Consumer surplus after dropping vehicles (millions):	Change:
\$35,902	\$36,050	\$149 (0.4%)

Table 7: Change in consumer surplus following the removal of four Chrysler models from the vehicle choice set by the merged firm. The measured change is positive, but close to 0.

The change is positive, which should not be possible, but it is only \$149 million, out of a market where the two firms combined are simulated to provide \$35.9 billion in consumer surplus. A similarly proportioned gain to consumer surplus was computed under multiple combinations of demand curve points being used for the approximation. This suggests that however much consumer surplus is harmed by the loss of these vehicles, the loss is so small that it disappears in the measurement errors of the linear approximations of the demand curves.

V. Conclusion

5.1 Discussion and policy implications

The merged firm's small revenue losses from the discontinuation of four of Chrysler's models suggest that a combined GM-Chrysler would likely cut several vehicles. This vehicle choice set does not even consider the even-lower selling Chrysler and GM models, which can reasonably be assumed to exhibit even higher fixed-cost/annual revenue ratios. As noted in the results, it seems the low drop in revenues is driven by consumers frequently substituting into one of the merged firm's remaining vehicles. However, certain high-revenue, iconic Chrysler brands such as the Dodge Ram

and Jeep Wrangler seem likely to survive. These findings bear out my hypotheses about product line consolidation.

On the other hand, the manufacturer loyalty can apparently only be relied on to a degree, because the 5% price hike designed to capture some rents from that loyalty turned out to cause heavy losses in revenue. This important finding suggests that on the whole, strong competition will remain in every vehicle segment, disciplining the merged firm by leaving it little room to raise prices. This analysis does not consider the prospect of a movement upward in every firm's equilibrium pricing decisions, but this result suggests that another firm could gain quite a bit by undercutting GM-Chrysler's prices near where they are now, so I would not expect the equilibrium prices to rise very much.

Finally, consumer welfare would only be infinitesimally affected, according to my simulations. This suggests that from a policy perspective, the downsides to a merger would not be too great. The substitution patterns evident in the simulated demand model here reinforce the arguments of those pressing for a merger and streamlining of the two firms: they could save significantly on costs by cutting brands and re-negotiating labor contracts, they have many redundant models among which many buyers are nearly indifferent (hence the negligible lost consumer surplus when they are forced to switch), and other brands, particularly a resilient Ford as well as the steadily growing imports, are a powerful disciplining force against any anti-competitive effects of the merger.

Works Cited

- Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63, no. 4: 841-890.
- Bresnahan, Timothy F. 1987. "Competition and Collusion in the American Automobile Industry: The 1955 Price War." *Journal of Industrial Economics* 35, no. 4: 457-482.
- Goldberg, Pinelopi Koujianou. 1995. "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry." *Econometrica* 63, no. 4: 891-951.
- Nevo, Aviv. 2000. "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry." *RAND Journal of Economics* 31, no. 3: 395-421.
- Petrin, Amil. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy* 110, no. 4: 705-729.
- Train, Kenneth. *Discrete Choice Methods with Simulation*. Cambridge University Press, 2003.
- Train, Kenneth E., and Clifford Winston. 2007. "Vehicle Choice Behavior and the Declining Market Share of U.S. Automakers." *International Economic Review* 48, no. 4: 1469-1496.
- Werden, Gregory J., Luke M. Froeb, and Timothy J. Tardiff. 1996. "The Use of the Logit Model in Applied Industrial Organization." *International Journal of the Economics of Business* 3, no. 1: 83-105.